



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/812,109	03/29/2004	James A. Mott	SUN-P8922	1538
57960 7590 02/19/2008 SUN MICROSYSTEMS INC. C/O PARK, VAUGHAN & FLEMING LLP 2820 FIFTH STREET DAVIS, CA 95618-7759			EXAMINER WONG, WARNER	
			ART UNIT 2616	PAPER NUMBER
			MAIL DATE 02/19/2008	DELIVERY MODE PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No. 10/812,109	Applicant(s) MOTT, JAMES A.	
	Examiner Warner Wong	Art Unit 2616	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 29 March 2004.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-14 and 30-43 is/are pending in the application.
- 4a) Of the above claim(s) 15-29 and 44-55 is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-14 and 30-43 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____ |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | 5) <input type="checkbox"/> Notice of Informal Patent Application |
| 3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date _____ | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

Election/Restrictions

1. Restriction to one of the following inventions is required under 35 U.S.C. 121:
 - I. Claims 1-10, 11-14, 30-43, drawn to flow control, classified in class 370, subclass 235.
 - II. Claims 15-25, 44-48, drawn to flow control, classified in class 370, subclass 235.
 - III. Claims 26-29, drawn to flow control, classified in class 370, subclass 235.
 - IV. Claims 49-55, drawn to determination of data admission, classified in class 370, subclass 232.
2. Inventions I and II are directed to related distinct processes. The related inventions are distinct if the (1) the inventions as claimed are either not capable of use together or can have a materially different design, mode of operation, function, or effect; (2) the inventions do not overlap in scope, i.e., are mutually exclusive; and (3) the inventions as claimed are not obvious variants. See MPEP § 806.05(j). In the instant case, the inventions as claimed describes different mode of operation: one invention is based on flow control requiring both queue pairs and virtual lanes and the other is based on flow control requiring only queue pairs. Furthermore, the inventions as claimed do not encompass overlapping subject matter and there is nothing of record to show them to be obvious variants.

3. Inventions I and III are directed to related distinct processes. The related inventions are distinct if the (1) the inventions as claimed are either not capable of use together or can have a materially different design, mode of operation, function, or effect; (2) the inventions do not overlap in scope, i.e., are mutually exclusive; and (3) the inventions as claimed are not obvious variants. See MPEP § 806.05(j). In the instant case, the inventions as claimed describes different mode of operation: one invention is based on flow control requiring both queue pairs and virtual lanes and the other is based on flow control requiring only virtual lanes. Furthermore, the inventions as claimed do not encompass overlapping subject matter and there is nothing of record to show them to be obvious variants.

4. Inventions I and IV are directed to related distinct processes. The related inventions are distinct if the (1) the inventions as claimed are either not capable of use together or can have a materially different design, mode of operation, function, or effect; (2) the inventions do not overlap in scope, i.e., are mutually exclusive; and (3) the inventions as claimed are not obvious variants. See MPEP § 806.05(j). In the instant case, the inventions as claimed describe different functions for remote nodes: one invention is about flow control and the other is about data admission. Furthermore, the inventions as claimed do not encompass overlapping subject matter and there is nothing of record to show them to be obvious variants.

5. Inventions II and III are directed to related distinct processes. The related inventions are distinct if the (1) the inventions as claimed are either not capable of use together or can have a materially different design, mode of operation, function, or effect; (2) the inventions do not overlap in scope, i.e., are mutually exclusive; and (3) the inventions as claimed are not obvious variants. See MPEP § 806.05(j). In the instant case, the inventions as claimed describes different mode of operation: one invention is based on flow control requiring only queue pairs and the other is based on flow control requiring only virtual lanes. Furthermore, the inventions as claimed do not encompass overlapping subject matter and there is nothing of record to show them to be obvious variants.

6. Inventions II and IV are directed to related distinct processes. The related inventions are distinct if the (1) the inventions as claimed are either not capable of use together or can have a materially different design, mode of operation, function, or effect; (2) the inventions do not overlap in scope, i.e., are mutually exclusive; and (3) the inventions as claimed are not obvious variants. See MPEP § 806.05(j). In the instant case, the inventions as claimed describe different functions for remote nodes: one invention is about flow control and the other is about data admission. Furthermore, the inventions as claimed do not encompass overlapping subject matter and there is nothing of record to show them to be obvious variants.

7. Inventions III and IV are directed to related distinct processes. The related inventions are distinct if the (1) the inventions as claimed are either not capable of use together or can have a materially different design, mode of operation, function, or effect; (2) the inventions do not overlap in scope, i.e., are mutually exclusive; and (3) the inventions as claimed are not obvious variants. See MPEP § 806.05(j). In the instant case, the inventions as claimed describe different functions for remote nodes: one invention is about flow control and the other is about data admission. Furthermore, the inventions as claimed do not encompass overlapping subject matter and there is nothing of record to show them to be obvious variants.

Because these inventions are independent or distinct for the reasons given above and there would be a serious burden on the examiner if restriction is not required because the inventions require a different field of search (see MPEP § 808.02), restriction for examination purposes as indicated is proper.

During a telephone conversation with Andrew Crain on 1/13/08, a provisional election was made with traverse to prosecute the invention of I, claims 1-13 and 32-34. Affirmation of this election must be made by applicant in replying to this Office action. Claims 14-18, 27-30, 19-26 and 31 are withdrawn from further consideration by the examiner, 37 CFR 1.142(b), as being drawn to a non-elected invention.

Claim Rejections - 35 USC § 103

8. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claim 1 is rejected under 35 U.S.C. 103(a) as being unpatentable over Hendel (US 2004/0120332) In view of Bloch (US 2001/0043564).

Regarding claim 1, Hendel describes a Method of flow controlling Infiniband receive traffic, comprising:

maintaining a single memory structure for queuing Infiniband traffic received via multiple virtual lanes and multiple queue pairs (para. 6 & 27-28, sharing the memory resource for the input traffic to the queue pair with corresponding virtual lane);

identifying a first packet payload received via a first virtual lane and a first queue pair (para. 38-39, send queues receives messages comprising encapsulated data);

determining whether the first payload can be stored in the memory structure without exceeding a portion of the memory structure allocated to the first queue pair (fig. 6, step 602 & 86, determine if available resources > 0 are available for the WQE (first queue pair, while having a minimum guarantee allocation for the queue (para. 6)));

if storing the first payload can be stored in the memory structure would exceed said portion of the memory structure allocated to the first queue pair, determining whether the first queue pair is enabled to use a shared portion of the memory structure

to store payloads of packets received via the first queue pair (fig. 6, step 626 & para. 100, determining if available resources < reserves (shared portion)).

Hendel fails to suggest:

determining whether the first payload can be stored in the memory structure without exceeding a portion of the memory structure allocated to the first virtual lane (fig. 4 & para. 48, determining if received packet X can be stored in the buffer).

Bloch describes:

determining whether the first payload can be stored in the memory structure without exceeding a portion of the memory structure allocated to the first virtual lane (fig. 4 & para. 48, determining if received packet X can be stored in the buffer).

It would have been obvious to one with ordinary skill in the art at the time of invention by applicant to incorporate the determination step of whether the data/payload can be stored in the memory structure without exceeding the VL allocation as per Bloch for the Infiniband memory allocation of Hendel.

The motivation for combining the teachings is that it provides an improved method for flow control in a packet switching fabric (Bloch, para. 6).

Regarding claim 2, Hendel describes:

allocating a portion of the memory structure to each of the multiple queue pairs (fig. 6, step 626 & para. 100-101, each send queue can accept a new buffer).

Hendel fails to describe:

allocating a portion of the memory structure to each of the multiple virtual lanes

Bloch describes:

allocating a portion of the memory structure to each of the multiple virtual lanes (para. 52).

Regarding claim 3, Bloch and Hendel combined describe:

memory structure comprises a set of linked lists of memory structure buffers, including one linked list for each of the multiple pairs that are active (para. 39, each queue/sub-queue is maintained in a form of linked list).

Regarding claim 4, Bloch and Hendel combined describe:

dropping the first payload if the first payload cannot be stored in the memory structure without exceeding the portion of the memory structure allocated to the first virtual lane (fig. 4, step of "dumping arriving packet").

Regarding claim 5, Block further suggests:

Issuing a Retry, Not Ready, Negative Ack (RNR-NAK) if:

The first payload cannot be store in the memory structure without exceeding a portion of the memory structure allocated to the first queue pair, and the first queue pair is not enabled to use the shared portion of the memory structure (abstract & fig. 6, step 628, a "no" response (Not Ready/NAK) for a current queue where "current > minimum" if available and reserve memory don't have enough resources, thus not allowing (enabling) allocation).

Regarding claim 6, Block further suggests:

The first payload cannot be store in the memory structure without exceeding a portion of the memory structure allocated to the first queue pair, and the first queue pair is enabled to use the shared portion of the memory structure and the shared portion of

the memory structure is full (abstract & fig. 6, step 628, a "no" response (Not Ready/NAK) for a current queue where "current > minimum" if available and reserve memory don't have enough resources).

Regarding claim 7, Hendel and Bloch combined further suggests:

defining one or more dedicated thresholds in the portion of the memory structure allocated to the first queue pair, and for each of said dedicated thresholds, identifying a number of message credits the queue (Bloch, para. 13, assigning maximum limits in form of transmission credits).

Regarding claim 9, Hendel further suggests:

receiving a request on a second queue pair to perform an RDMA (Remote Direct Memory Access) Read operation, and based on an amount of data expected to be received via RDMA Read operation, reserving a sufficient number of buffers in the memory structure (para. 38, the dynamic allocating of shared pool of buffers are based on RDMA read request).

Regarding claim 10, Hendel further suggests:

in the single memory structure, reassembling the queued Infiniband traffic into outbound communications (fig. 6, step 612, allocated WQE (where client data resides) to send queue (outbound communications);

receiving a payload on an idle queue pair, wherein a queue pair is idle if no traffic from the queue pair is stored in the single memory structure (fig. 2 & para. 31, receiving client data on an [empty] (idle) send queue (queue pair));

only queuing the payload in the single memory structure if sufficient space in the single memory structure is reserved for completing reassembly of outbound communication on each non-idle queue pair (fig. 6, step 628 & para. 102, deploying certain current send queue data for which $\text{current} < \text{minimum}$) .

9. Claim 8 is rejected under 35 U.S.C. 103(a) as being unpatentable over Hendel in view of Bloch as applied to claim 1 above, and further in view of Pekkala (US 2002/0085493).

Regarding claim 8, Hendel and Bloch combined further describes:

defining one or more shared thresholds in the shared portion of the memory structure (Bloch, par. 27, limits in form of credits for the receive buffer);

but fail to describe:

for each shared thresholds, identifying a number of message credits the queue pair may advertise when the amount of shared portion used by the multiple queue pairs exceeds said shared threshold.

Pekkala describes:

for each shared threshold, identifying a number of message credits the queue pair may advertise when the amount of shared portion used by the multiple queue pairs exceeds said shared threshold (para. 19, if free space in memory drops below a predetermined threshold, flow control credits are advertised).

It would have been obvious to one with ordinary skill in the art at the time of invention by applicant to describe using thresholds for advertising message credits under a threshold/limit as in Pekkala for the Infiniband system of Hendel and Bloch.

The motivation for combining the teachings is that it provides an Infiniband buffering scheme which maintains an acceptable level of performance in a realistic manufacturable manner (Bloch, para. 14)

Regarding claim 11, it is a computer readable medium storing executable instructions claim comprising limitation of the method of claim 1. Hence, it is rejected under the same rationale.

Claims 12, 13 and 14 comprise limitations which correspond to claims 8, 5 and 6 respectively. Hence, it is rejected under the same rationale.

Claim Rejections - 35 USC § 102

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

10. Claims 30-31 are rejected under 35 U.S.C. 102(e) as being anticipated by Hendel.

Regarding claim 30, Hendel describes a method comprising:

maintaining a single memory structure for reassembling InfiniBand traffic received via multiple virtual lanes and multiple queue pairs (abstract, using shared pool

of buffers to arrange transmission of received data using Infiniband virtual lanes and queue pairs, para. 27-28);

identifying a first packet payload received via a first queue pair that is idle, wherein the first queue pair is considered idle if no traffic from the first queue pair is stored in said single memory structure (fig. 2 & para. 31, receiving client data on an [empty] (idle) send queue (queue pair));

for each other queue pair for which traffic from said queue pair is stored in said single memory structure, determining whether sufficient space in the single memory structure is reserved for reassembling said traffic, and storing the first packet payload in said single memory structure only if sufficient space (fig. 6, step 628 & para. 102, deploying certain current send queue data for which $\text{current} < \text{minimum}$).

Regarding claim 31, Hendel further describes:

identifying an amount of space in said single memory structure reserved for said other queue pair (fig. 5, step 610, finding unallocated WQE buffer);

comparing said amount of reserved space to an amount of space expected to be needed to complete reassembly of said traffic from said other queue pair (fig. 6, step 612, allocated WQE (where client data resides) to send queue (outbound communications)).

11. Claim 32-37 are rejected under 35 U.S.C. 103(a) as being unpatentable over Hendel and Pekkala.

Regarding claim 32, Hendel describes an apparatus for flow controlling received Infiniband traffic, comprising:

a single memory structure configured to queue payloads of InfiniBand traffic received via multiple virtual lanes and multiple queue pairs (para. 6, 27-28, shared pool of buffers used for queuing packets in queue pairs and virtual lanes);

a resource manager configured to manage the memory structure (para. 8, 104 and fig. 7; queue's context 608 & controller 602 control the shared pool of buffers for queuing);

Hendel fails to describe:

a first module configured to facilitate the advertisement of virtual lane credits.

Pekkala describes:

An apparatus configured to facilitate the advertisement of virtual lane credits (abstract, credit flow control packets for each of the VLs).

It would have been obvious to one with ordinary skill in the art at the time of invention by applicant to incorporate the controlling apparatus which advertises message credits for virtual lanes as in Pekkala for the Infiniband system of Hendel.

The motivation for combining the teachings is that it provides an Infiniband buffering scheme which maintains an acceptable level of performance in a realistic manufacturable manner (Bloch, para. 14)

Hendel describes a system (fig. 7) which modules supporting the InfiniBand queue pairs. Hence, it inherently describes a second module configured to facilitate the advertisement of queue pair credits (the examiner cites but not used as a reference

Kagan, US 2004/0017819, para. 9, where the InfiniBand specification already provides a flow control mechanism where each credit represents 1 WQE corresponding to a QP).

Regarding claim 33, Hendel further describes:

single memory structure comprises multiple linked lists of memory structure buffers, including one linked list for each of the multiple queue pairs that is active (para. 57, allocated list 404 is a linked lists) of buffers allocated to different send queues).

Regarding claim 34, Hendel further suggests:

the first module comprises an InfiniBand link core (fig. 7, each of the controlling apparatus module is consider essential (core) to the Infiniband system).

Regarding claim 35, Hendel further suggests:

second module comprises an acknowledgement generator configured to generate transport layer acknowledgements (fig. 7 & para. 43, mux 304 forwards (generates) ACKs to other endpoints).

Regarding claim 36, Hendel further describe:

a processor interface configured to facilitate the programming of operating parameters associated with the multiple virtual lanes and multiple queue pairs (fig. 7, mux 324 is an interface to context 608 & controller 602 (processor) which handles VLs and queue pairs).

Regarding claim 37, Hendel and Pekkala combined suggest:

a first memory configured to store one or more parameters associated with operation of a first virtual lane (Pekkala, abstract, storing the quantity of buffer resources (parameter) for each virtual lane).

12. Claim 38-41 are rejected under 35 U.S.C. 103(a) as being unpatentable over Hendel in view of Pekkala as applied to claim 37 above, and further in view of Bloch.

Regarding claim 38, Hendel and Pekkala combined suggest:

a count of the number of memory structure buffers currently used to store payload of packets received via the first virtual lane (Pekkala, abstract, storing the quantity of buffer resources for each virtual lane).

Pekkala also describes a threshold in accounting in relation to credits (abstract, shutdown latency threshold), but fails to describe:

a first packet is dropped if storing the payload of the first packet would cause said count to exceed said threshold.

Bloch describes:

a first packet is dropped if storing the payload of the first packet would cause said count to exceed said threshold (fig. 4 & para. 48, $x \leq \text{free credits}$, dumping arriving packet).

It would have been obvious to one with ordinary skill in the art at the time of invention by applicant to incorporate the determination step of dropping the packet regarding a threshold as per Bloch for the flow control of Infiniband memory allocation of Hendel.

The motivation for combining the teachings is that it provides an improved method for flow control in a packet switching fabric (Bloch, para. 6).

Regarding claim 39, Hendel further describes:

a second memory configured to store, for each of the multiple queue pairs that is active, one or more parameters associated with operation of said queue pair (para. 6, (second) memory to store current amount of resource allocated to the queue, minimum guarantees,

Regarding claim 40, Hendel further describes:

a maximum number of memory structure buffers dedicated to storing payloads of packets received via said queue pair (fig. 4 & para. 6, maximum allocation to the queue regardless of other queues' allocations).

an indicator configured to indicate whether said queue pair is enabled to use a set of shared memory structure buffers (fig. 6, step 624, comparison (indicator) whether send queue can use all of reserve memory).

Hendel fails to describe: a maximum number of message credits advertisable by said queue pair

Bloch describes: a maximum number of message credits advertisable by said queue pair (para. 42).

Regarding claim 41, Hendel and Pekkala combined further suggest:

one or more dedicated thresholds, wherein each said dedicated threshold identifies a subset of said maximum number of memory structure buffers (Pekkala, para. 22, predetermined threshold of buffering resources for each port);

for each said dedicated threshold, a number of message credits advertisable by said queue pair when said queue pair uses said subset of said maximum number of

memory structure buffers (Pekkala, para. 22-23, advertisement of amount of credits when above threshold).

Allowable Subject Matter

13. **Claims 42-43** are objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims.

The following is a statement of reasons for the indication of allowable subject matter:

The prior art fails to describe an Infiniband system further comprising:

“ a number of shared memory structure buffers in said set of shared memory structure buffers, wherein said shared memory structure buffers are available for use by said queue pair to store payloads of packets received via said queue pair if:

Said queue pair has used said maximum number of memory structure buffers;
said indicator indicates that said queue pair is enabled to use said set of shared memory structure buffers;

a maximum number of message credits advertisable by said queue pair when said queue pair starts using shared memory structure buffers.”

Conclusion

14. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure: Pettey (US 2007/0025354) describing method for shared I/O in

load/store fabric, Arndt (US 2005/0018669) describing Infiniband subnet management queue pair emulation on a single physical port, Gupta (US 2005/0089033) describing an Infiniband channel adaptor, Rojas (US 2003/0223416) describing dynamic reallocation of virtual lane buffer space in an Infiniband switch, Gil (US 6,904,507) describing buffer management architecture for Infiniband subnetwork, Kagan (US 2004/0017819) describing an Infiniband receive queue descriptor pool, Susnow (US 2002/0159385) describing a link level packet flow control mechanism and Pekkala (US 2005/0172195) describing the transaction buffering within Infiniband Device.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Warner Wong whose telephone number is 571-272-8197. The examiner can normally be reached on 6:30AM - 3:00PM, M-F.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Kwang Yao can be reached on 571-272-3182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Warner Wong
Examiner
Art Unit 2616

WW

KWANG BIN YAO
SUPERVISORY PATENT EXAMINER

